

Signal Enhancement as Minimization of Relevant Information Loss

Bernhard C. Geiger*, Gernot Kubin*

*Signal Processing and Speech Communication Laboratory, Graz University of Technology, Austria
{geiger, gernot.kubin}@tugraz.at

Abstract—We introduce the notion of relevant information loss for the purpose of casting the signal enhancement problem in information-theoretic terms. We show that many algorithms from machine learning can be reformulated using relevant information loss, which allows their application to the aforementioned problem. As a particular example we analyze principle component analysis for dimensionality reduction, discuss its optimality, and show that the relevant information loss can indeed vanish if the relevant information is concentrated on a lower-dimensional subspace of the input space.

I. INTRODUCTION

It is a widely known fact from information theory, that processing a random variable, or a signal, cannot increase the amount of information it represents. In fact, as the data processing inequality states, information can only be lost by passing a signal through a deterministic system. This information loss can – very loosely speaking – be interpreted as the difference between the information at the input and the information at the output of the system. In case of continuous-valued random variables (which contain an infinite amount of information) the authors developed a theory quantifying the information loss both in absolute and relative terms [1], [2].

We now look at the data processing inequality (DPI) from a different point of view by asking the following question: How can we justify *signal processing*, knowing that it can only *reduce* information? Whenever technical systems prepare a physical, information-carrying signal for human perception, processing occurs and, presumably, a significant amount of information is lost. The only justification for this approach is that the information which is lost is actually not the information we are interested in, but rather some nuisance. We hope to preserve all the information *relevant* to us while minimizing the nuisance. The problem of signal enhancement, as we will argue later, is an optimization problem which can be cast in exactly these terms. Aside from that, each block in the signal processing chain has the goal of representing the relevant information such that as little as possible is lost in the subsequent blocks; as in the sense of [3], preprocessing can improve performance, while postprocessing cannot – we cannot recover lost information, but we can prevent losing it. It is our goal to make these statements precise by supporting them with a mathematical framework of *relevant information loss*.

Relevant information and its counterpart, relevant information loss, are concepts not altogether new. Indeed, the

notion of relevance has a long history in machine learning and neural networks: Plumbley, essentially using the same definition as we do, explicitly used relevant information loss in analyzing the properties of principle component analysis [4], [5]. The information bottleneck (IB) method and its variants directly maximize relevant information while minimizing the complexity of its signal representation [6]. Relevance w.r.t. a specific goal was the motivation for defining an information processing measure for neural networks in [7], [8]. Finally, the principle of relevant information, being structurally similar to the information bottleneck formulation, builds on minimizing the relative entropy between the relevant random variable and its representation [9].

Conversely, in signal processing the notion of relevant information has not gained foothold yet. In this paper we thus formulate a definition of relevant information loss (Section II) following the ideas of [4] and discuss its properties from the view-point of system theory (Section III). On the basis of these properties we then discuss the problem of signal enhancement in Section IV and investigate the connections to the machine learning literature. As a special example of signal enhancement we analyze principle component analysis (Section V), generalizing the results of [4]. Finally, we give a brief example of the application of our results to a simple digital communication system (Section VI).

II. A DEFINITION OF RELEVANT INFORMATION LOSS

We start with recalling the definition given in [1], where the information loss induced by transforming a (possibly multidimensional) random variable (RV) X to another RV Y by a static function $g: \mathcal{X} \rightarrow \mathcal{Y}$, $\mathcal{X}, \mathcal{Y} \subseteq \mathbb{R}^N$ was given as

$$L(X \rightarrow Y) = \sup_{\mathcal{P}} \left(I(\hat{X}; X) - I(\hat{X}; Y) \right) = H(X|Y) \quad (1)$$

where the supremum is over all partitions \mathcal{P} of the sample space \mathcal{X} of X , and where \hat{X} is obtained by quantizing X accordingly (see Fig. 1(a)). This definition of absolute information loss is accompanied by an expression for the relative information loss in [2].

As the examples in [2] show, both the absolute and the relative definition of information loss have their shortcomings, especially when it comes to systems g used for signal enhancement: Since the expressions only consider the RV X at the input of the system, they do not take into account that not all of the information contained in X is *relevant*, often leading to

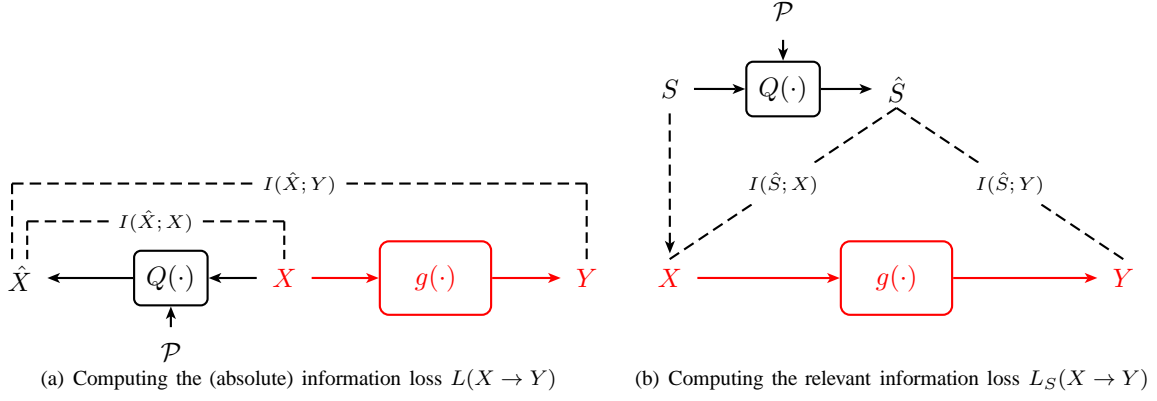


Fig. 1. Model for computing the information loss of a memoryless input-output system g . Q is a quantizer with partition \mathcal{P} .

counter-intuitive results. The main contribution of this work lies thus in analyzing the implications of the following

Definition 1 (Relevant Information Loss). Let X be an RV on the sample space \mathcal{X} , and let Y be obtained by transforming X with a static function g . Let S be another RV on the sample space \mathcal{S} representing *relevant information*. The information loss *relevant w.r.t. S* is defined as

$$L_S(X \rightarrow Y) = \sup_{\mathcal{P}} \left(I(\hat{S}; X) - I(\hat{S}; Y) \right) = I(X; S|Y). \quad (2)$$

where the supremum is over all partitions of the sample space \mathcal{S} , and where \hat{S} is obtained by quantizing S accordingly.

The information loss relevant w.r.t. an RV S is thus the difference of mutual informations between S and the input and output of the system (see Fig. 1(b)). Equivalently, the relevant information loss is exactly the information that X contains about S which is not contained in Y . Due to the data processing inequality (e.g., [10]) it is a non-negative quantity, i.e., a deterministic system cannot increase the amount of available relevant information.

As we have already pointed out in the introduction, this definition of relevant information loss is not altogether new: Plumbley already introduced this quantity (named $\Delta I_S(X; Y)$), and omitting the supremum assuming finite mutual information between S and X in the context of unsupervised learning in neural networks [4]. The rationale for this new measure was to circumvent some shortcomings of Linsker's principle of information maximization [11]: While infomax works well for Gaussian RVs and linear systems, applying the same algorithms to non-Gaussian data just maximizes an upper bound on the information. Plumbley's information loss, conversely, also yields closed form solutions for Gaussian data, but in addition to that one can derive upper bounds on the relevant information loss whenever the data is non-Gaussian. Thus, minimizing an upper bound on the information loss can be assumed to be more promising than maximizing an upper bound on the transferred information [4].

III. PROPERTIES OF RELEVANT INFORMATION LOSS

We now analyze the elementary properties of relevant information loss:

Proposition 1 (Elementary Properties). *The relevant information loss from Definition 1 satisfies the following properties:*

- 1) $L_S(X \rightarrow Y) \leq I(S; X) \leq H(S)$
- 2) $L_S(X \rightarrow Y) \geq L_Y(X \rightarrow Y) = 0$
- 3) $L_S(X \rightarrow Y) \leq L_X(X \rightarrow Y) = L(X \rightarrow Y)$, with equality if X is a function of S .
- 4) $L_S(X \rightarrow Y) = H(S|Y)$ if S is a function of X .

Proof: The first property results immediately from the definition, while the second property is due to the fact that Y is a function of X . The third property results from making X the relevant RV, thus making (2) equal to (1). Note that, with (1),

$$L_S(X \rightarrow Y) = I(S; X|Y) = H(X|Y) - H(X|Y, S) \quad (3)$$

$$\leq L(X \rightarrow Y) \quad (4)$$

with equality if X is a function of S . The last property, for $S = f(X)$, follows by expanding $I(X; S|Y)$ as $H(S|Y) - H(S|Y, X) = H(S|Y)$. ■

The third property is of particular interest. Essentially, it states that the relevant information loss cannot exceed the total information loss, so upper bounds (e.g., those presented in [1]) for the latter can be used for the former as well. In addition to that, it suggests an alternative way to define $L(X \rightarrow Y)$: Namely as

$$L(X \rightarrow Y) = \sup_S (I(S; X) - I(S; Y)) \quad (5)$$

where the supremum is over all RVs on the sample space \mathcal{X} .

Since by Definition 1 relevant information loss is represented by a conditional mutual information, it inherits all of its properties. In particular, as we show next, a data processing inequality (DPI) holds:

Proposition 2 (Data Processing Inequality). *Let $V - W - X - Y$ be a Markov chain. Then,*

$$L_W(X \rightarrow Y) \geq L_V(X \rightarrow Y). \quad (6)$$




Fig. 2. Cascade of two systems: The relevant information loss of the cascade is the sum of relevant information losses of the constituent systems.

Proof: See Appendix. ■

The particular usefulness of the DPI (especially in the context of our work) relies on the fact that both $S - X - g(X)$ and $f(S) - S - X$ are Markov chains. Comparing this to Proposition 2 one is tempted to believe that the direction of the inequality depends on the fact whether $f(S) - S - X - Y$ or $S - f(S) - X - Y$ is a Markov chain. The following Corollary resolves this complication.

Corollary 1. *Let f be a measurable function defined on the sample space of S . Then,*

$$L_S(X \rightarrow Y) \geq L_{f(S)}(X \rightarrow Y) \quad (7)$$

with equality if $S - f(S) - X - Y$ is a Markov chain.

Proof: See Appendix. ■

Before proceeding, we note that the third property of Proposition 1 is an immediate consequence of this corollary.

We next show that, inherited from the properties of mutual information, also the relevant information loss obeys a chain rule:

Proposition 3 (Chain Rule of Information Loss). *The information loss $L_{S_1^n}(X \rightarrow Y)$ w.r.t. a collection $S_1^n = \{S_1, \dots, S_n\}$ of relevant RVs induced by a function g satisfies*

$$L_{S_1^n}(X \rightarrow Y) = \sum_{i=1}^n L_{S_i|S_1^{i-1}}(X \rightarrow Y). \quad (8)$$

The proof follows immediately from the chain rule of (conditional) information and is thus omitted. However, this chain rule justifies our intuitive understanding of the nature of information loss. We thus emphasize

Corollary 2. *The information loss $L(X \rightarrow Y)$ induced by a function g can be split into relevant (w.r.t. S) and irrelevant information loss:*

$$L(X \rightarrow Y) = L_S(X \rightarrow Y) + L_{X|S}(X \rightarrow Y) \quad (9)$$

Proof: The proof follows from Proposition 3 and from the fact that since S and Y are conditionally independent given X we have $L_{XS}(X \rightarrow Y) = L(X \rightarrow Y)$. ■

Note that even in a simple scenario with additive noise, i.e., $X = S + N$, it is not straightforward to just identify the noise N with the irrelevant information. We will show this in one of the examples in Section VI.

A final property we want to present concerns the cascade of systems (see Fig. 2). Assume, for the moment, that the output




Fig. 3. The problem of signal enhancement: Given a (relevant) information signal S and its observation X , how shall we design a system such that its output, Y , represents S as good as possible?

Y of the first system g is used as the input to a second system h , which responds with the RV Z . In [12] it was shown that the information loss of this cascade is additive, i.e.,

$$L(X \rightarrow Z) = L(X \rightarrow Y) + L(Y \rightarrow Z). \quad (10)$$

The same holds for the relevant information loss, as Plumbley showed in [4]. We thus state

Proposition 4 (Information Loss in a Cascade is Additive, [4]). *Let Y be the RV obtained by transforming X with a static function g , and let Z be obtained by transforming Y with h . Then, the information loss relevant w.r.t. S is given as*

$$L_S(X \rightarrow Z) = L_S(X \rightarrow Y) + L_S(Y \rightarrow Z). \quad (11)$$

Proof: See Appendix. ■

IV. SIGNAL ENHANCEMENT, RELEVANT INFORMATION LOSS, AND IB WITH SIDE INFORMATION

We now turn our attention to the signal enhancement problem. Often, the information one wants to retrieve is not directly available, but only through some corrupted observation: Information-carrying signals are superimposed by noise, distorted through nonlinear systems, and affected by time-varying, dynamic effects. It is essentially the goal of the signal processing engineer to mitigate all these adverse effects; to improve the quality of the observation such that as much information as possible can be retrieved from it with little effort.

Looking at this task from the perspective of information theory, we know that, as the data processing inequality dictates, signal enhancement does *not* mean that one increases the amount of information in the observation. At best, one can build a system which *preserves* as much information as possible (see Fig. 3). Conversely, noise, distortion, and other *irrelevant* components of the observation should be removed such that information retrieval can be done easily.

This is where our notion of relevant information loss comes in: If we let S be the information carrying signal and X its corrupted observation, our goal is to find a function g such that the relevant information loss is minimized while simultaneously maximizing the irrelevant information loss. One may cast the resulting optimization problem as follows:

$$\begin{aligned} \max_g \quad & L_{X|S}(X \rightarrow Y) \\ \text{s.t.} \quad & L_S(X \rightarrow Y) \leq C \end{aligned} \quad (12)$$

where C is some constant¹. In other words, signal enhancement is the *maximization of irrelevant information loss*.

Although in signal processing these information-theoretic considerations are only slowly gaining momentum, they are quite widespread in machine learning and neural networks, e.g., [5], [9]. In particular, a variational formulation of maximizing relevant information is the basis of the information bottleneck method (IB) [6]:

$$\min_{p(y|x)} I(Y; X) - \beta I(S; Y) \quad (13)$$

where the minimization is performed over all relations between the (discrete) RVs Y and X and where β is a design parameter trading compression for preservation of relevant information, $I(S; Y)$. While in principle the relation $p(y|x)$ can be stochastic, in many cases the algorithm is used for (hard) clustering using a deterministic function. We now try to express (13) in terms of relevant and irrelevant information loss: Using $L_S(X \rightarrow Y) = I(S; X) - I(S; Y)$ we obtain

$$I(S; Y) = I(S; X) - L_S(X \rightarrow Y) \quad (14)$$

where the first term is obviously independent of $p(y|x)$. Restricting ourselves to the clustering problem, i.e., to deterministic functions $Y = g(X)$, and letting g° be the optimal solution, we get

$$g^\circ = \arg \min_g I(Y; X) - \beta I(S; Y) \quad (15)$$

$$= \arg \min_g -H(X) + I(Y; X) - \beta I(S; X) + \beta L_S(X \rightarrow Y) \quad (16)$$

$$= \arg \min_g -L(X \rightarrow Y) + \beta L_S(X \rightarrow Y) \quad (17)$$

$$= \arg \min_g -L_S(X \rightarrow Y) - L_{X|S}(X \rightarrow Y) + \beta L_S(X \rightarrow Y) \quad (18)$$

$$= \arg \min_g (\beta - 1)L_S(X \rightarrow Y) - L_{X|S}(X \rightarrow Y) \quad (19)$$

and thus, the optimization problem can be cast as

$$\min_g (\beta - 1)L_S(X \rightarrow Y) - L_{X|S}(X \rightarrow Y). \quad (20)$$

Note that for large β stronger emphasis is placed on minimizing relevant information loss, and a trivial solution to this problem is obtained by any bijective g .

In the case of discrete RVs this problem can be circumvented by another variant of the IB method: The agglomerative IB method proposed in [13] starts with $Y = X$ and iteratively merges elements of the state space \mathcal{Y} of Y such that the relevant information loss is minimized in each step. By stopping this algorithm as soon as during some merging step $L_S(X \rightarrow Y) > C$, at least a local optimum of our original signal enhancement problem (12) can be found.

A step further has been made by the authors of [14], who extended the IB method to incorporate knowledge about

irrelevant signal components. They accompany the relevant information S by an irrelevance variable \bar{S} , assuming conditional independence between S and \bar{S} given X , and minimize the following functional:

$$I(Y; X) - \beta [I(Y; S) - \gamma I(Y; \bar{S})] \quad (21)$$

At the same time, the entropy of the system output Y and the irrelevant information $I(Y; \bar{S})$ are minimized, and the relevant information $I(Y; S)$ is maximized. As in IB, β and γ are the weights for these three conflicting goals. We can again use our notions of relevant and, employing Corollary 2, irrelevant information loss to rewrite the optimization problem. If we restrict ourselves to clustering, one obtains

$$\min_g (\beta - 1)L_S(X \rightarrow Y) - (\beta\gamma + 1)L_{X|S}(X \rightarrow Y) \quad (22)$$

where again all constant terms have been dropped. Note that here $X|S$ takes the place of \bar{S} , automatically fulfilling the requirement of conditional independence in [14].

Unlike for IB, by employing side information it does make sense to let $\beta \rightarrow \infty$: Compression emerges naturally from maximizing the irrelevant information loss (due to $\gamma > 0$). This is exactly the approach that has been taken up by the authors of [7], [8], who introduced the following information processing measure for discrete RVs X , Y , and S :

$$\Delta P(X \rightarrow Y|S) = H(X|S) - H(Y|S) - \alpha (H(S|Y) - H(S|X)) \quad (23)$$

The authors argue that the first difference in (23) corresponds to complexity reduction (i.e., reduction of irrelevant information), while the second term accounts for loss of relevant information. Indeed, with

$$L_{X|S}(X \rightarrow Y) = I(X; X|Y, S) = H(X|Y, S) \quad (24)$$

$$= H(X|S) - H(Y|S) \quad (25)$$

we can rewrite $\Delta P(X \rightarrow Y|S)$ with our notions of relevant information loss and state a variational utility function (to be maximized) for our signal enhancement problem (12):

$$\Delta P(X \rightarrow Y|S) = L_{X|S}(X \rightarrow Y) - \alpha L_S(X \rightarrow Y) \quad (26)$$

Here, $\alpha > 0$ is a design parameter trading between the loss of relevant and irrelevant information.

Aside from the IB method and its variants widely used in machine learning, many other functions and algorithms inherently solve our problem of signal enhancement: Quantizers used in digital communication systems and regenerative repeaters have the goal to preserve the information-carrying (discrete) RV as much as possible while removing continuous-valued noise. Bandpass filters (although not yet directly tractable using our notions of relevant information loss) remove out-of-band noise while leaving the information signal unaltered. And finally, methods for dimensionality reduction (such as the principle component analysis) remove redundancy and irrelevance while trying to preserve the interesting part of the observation. In the next section we focus on exactly this class of methods.

¹In some cases it may be beneficial to minimize the relevant information loss subject to a minimum reduction of irrelevant information, or even use a variational formulation.

V. OPTIMALITY OF THE PCA IN TERMS OF RELEVANT INFORMATION LOSS

We now analyze the relevant information loss of the principal component analysis (PCA). In this regard, we initially follow the reasoning of Plumbly [4] who showed that the PCA in some cases minimizes the relevant information loss. We then generalize Plumbly's results and even show that this loss can vanish. To this end, we define the PCA as

$$\mathbf{Y}_M = \mathbf{g}(\mathbf{X}) = \mathbf{I}_M \mathbf{Y} = \mathbf{I}_M \mathbf{W}^T \mathbf{X} \quad (27)$$

where \mathbf{X} is an N -dimensional continuous-valued input RV, \mathbf{Y}_M an M -dimensional output RV (composed of the first M elements of \mathbf{Y}), and \mathbf{W} the matrix of eigenvectors of the covariance matrix of \mathbf{X} , $\mathbf{C}_X = \mathbb{E}\{\mathbf{X}\mathbf{X}^T\}$. Thus, $\det \mathbf{W} = 1$ and $\mathbf{W}^{-1} = \mathbf{W}^T$. Furthermore, let \mathbf{I}_M be an $(M \times N)$ -matrix with ones in the main diagonal. The PCA is here used for dimensionality reduction and, assuming that one has perfect knowledge of the rotation matrix \mathbf{W} , the *relative* information loss equals $\frac{N-M}{N}$, while the absolute information loss is infinite [2].

It is known that the PCA minimizes the mean squared error for a reconstruction $\mathbf{X}_M = \mathbf{W}\mathbf{I}_M^T \mathbf{Y}_M$ of \mathbf{X} . In addition to that, as Linsker pointed out in [11], given that \mathbf{X} is an observation of a Gaussian RV \mathbf{S} corrupted by Gaussian noise, the PCA maximizes the mutual information $I(\mathbf{S}; \mathbf{Y})$. For non-Gaussian \mathbf{S} (and Gaussian noise with spherical symmetry), the PCA not only provides an upper bound on the mutual information, but also an upper bound on the relevant information loss [4], [5].

We now want to generalize the results such that non-iid and non-Gaussian noise is taken into account as well. In particular, we show that the relevant information loss can indeed vanish under some circumstances. Let

$$\mathbf{X} = \mathbf{S} + \mathbf{N} \quad (28)$$

where \mathbf{S} and \mathbf{N} are the relevant information and the noise, respectively, with covariance matrices \mathbf{C}_S and \mathbf{C}_N . We further assume that \mathbf{S} and \mathbf{N} are independent and, as a consequence, $\mathbf{C}_X = \mathbf{C}_S + \mathbf{C}_N$. Let $\{\lambda_i\}$, $\{\mu_i\}$, and $\{\nu_i\}$ be the sets of eigenvalues of \mathbf{C}_X , \mathbf{C}_N , and \mathbf{C}_S , respectively. Let the eigenvalues be ordered descendingly, i.e.,

$$\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_N. \quad (29)$$

We now write

$$\mathbf{Y}_M = \mathbf{I}_M \mathbf{W}^T \mathbf{X} = \mathbf{I}_M \mathbf{W}^T \mathbf{S} + \mathbf{I}_M \mathbf{W}^T \mathbf{N} = \tilde{\mathbf{S}}_M + \tilde{\mathbf{N}}_M. \quad (30)$$

As mentioned before, \mathbf{Y}_M is composed of the first M elements of the vector \mathbf{Y} . Conversely, we let \mathbf{Y}_c , $\tilde{\mathbf{S}}_c$, and $\tilde{\mathbf{N}}_c$ denote the last $N - M$ elements of the corresponding vectors. If, as in our case, the orthogonal matrix \mathbf{W} performs the PCA, the covariance matrix of \mathbf{Y}_M is a diagonal matrix with the M largest eigenvalues of \mathbf{C}_X . We now present

Lemma 1. *For above signal model, the relevant information loss in the PCA is given by*

$$L_S(\mathbf{X} \rightarrow \mathbf{Y}_M) = h(\mathbf{Y}_c | \mathbf{Y}_M) - h(\tilde{\mathbf{N}}_c | \tilde{\mathbf{N}}_M). \quad (31)$$

Proof: See Appendix. ■

Before proceeding, we need the following

Definition 2. Let X and Y be continuous RVs with arbitrary continuous (joint) distribution, and let X_G and Y_G be Gaussian RVs with same (joint) first and second moments. We define the conditional divergence

$$J(X|Y) = D(X|Y || X_G|Y_G) = h(X_G|Y_G) - h(X|Y). \quad (32)$$

Clearly, this quantity inherits all its properties from the Kullback-Leibler divergence, e.g., non-negativity, and can be considered as a measure of Gaussianity. We want to point out that despite the similarity, this quantity is not to be confused with *negentropy*. For negentropy, instead of using a jointly Gaussian distribution for (X_G, Y_G) for computing $h(X_G|Y_G)$ one uses a Gaussian distribution for $X_G|Y = y$ to compute $h(X_G|Y = y)$ before taking the expectation w.r.t. the true marginal distribution of Y .

While most results about the optimality of the PCA are restricted to the Gaussian case, we now employ conditional divergence to generalize some of these results. In particular, we maintain

Theorem 1. *Let $\mathbf{X} = \mathbf{S} + \mathbf{N}$, with \mathbf{N} and \mathbf{S} independent, and let \mathbf{Y}_M be obtained by performing dimensionality-reducing PCA. If \mathbf{N} is iid (i.e., \mathbf{C}_N is a scaled identity matrix) and more Gaussian than \mathbf{Y} in the sense*

$$J(\tilde{\mathbf{N}}_c | \tilde{\mathbf{N}}_M) \leq J(\mathbf{Y}_c | \mathbf{Y}_M) \quad (33)$$

then the PCA minimizes the Gaussian upper bound on the relevant information loss $L_S(\mathbf{X} \rightarrow \mathbf{Y}_M)$.

Proof: See Appendix. ■

This theorem essentially generalizes the result by Plumbly [4], [5], who claimed that the PCA minimizes the Gaussian upper bound on the information loss for spherical Gaussian noise (i.e., for Gaussian \mathbf{N} with \mathbf{C}_N being a scaled identity matrix). In addition to that, it helps justifying our use of information loss instead of information transfer; an upper bound on the latter would not be useful in our signal enhancement problem.

As a next step, we restrict ourselves to relevant information which is concentrated on an $L \leq M$ -dimensional subspace, but drop the requirement that \mathbf{N} is iid. We still assume, however, that \mathbf{C}_N (and, thus, \mathbf{C}_X) is full rank. Note that due to these assumptions $\lambda_i > 0$ and $\mu_i > 0$ for all i , while $\nu_i = 0$ for $i > L$. We are now ready to state

Theorem 2 (Bounds for the PCA). *Assume that \mathbf{S} has covariance matrix \mathbf{C}_S with at most rank M , and assume that \mathbf{N} is independent of \mathbf{S} and has (full-rank) covariance matrix \mathbf{C}_N . Let further \mathbf{N} be more Gaussian than \mathbf{Y} in the sense*

$$J(\tilde{\mathbf{N}}_c | \tilde{\mathbf{N}}_M) \leq J(\mathbf{Y}_c | \mathbf{Y}_M) \quad (34)$$

where \mathbf{Y}_M is obtained by employing the PCA for dimensionality reduction. Then, the relevant information loss is bounded




Fig. 4. Digital Communication System: The input signal S is uniformly distributed on $\{-1, 1\}$, the channel adds uniform noise to the signal. The channel output, X , is quantized using a set of thresholds $\{\gamma_i\}$ to obtain the decision variable Y .

from above by

$$L_S(\mathbf{X} \rightarrow \mathbf{Y}_M) \leq \frac{1}{2} \ln \left(\prod_{i=M+1}^N \frac{\mu_1}{\mu_i} \right) \quad (35)$$

where $\{\mu_i\}$ is the set of decreasing eigenvalues of $\underline{\mathbf{C}}_N$ and where \ln is the natural logarithm.

Proof: See Appendix. ■

Note that this upper bound is non-negative, since – by assumption on the ordering of the eigenvalues – any term in the product cannot be smaller than one.

We finally show that there are cases where, despite the fact that we reduce the dimensionality of the data, we are still able to preserve all of the relevant information:

Corollary 3. Assume that \mathbf{S} has covariance matrix $\underline{\mathbf{C}}_S$ with at most rank M , and assume that \mathbf{N} is zero-mean Gaussian noise independent of \mathbf{S} with covariance matrix $\sigma_N^2 \mathbf{I}$. Let \mathbf{Y}_M be obtained by employing the PCA for dimensionality reduction. Then, the relevant information loss $L_S(\mathbf{X} \rightarrow \mathbf{Y}_M)$ vanishes.

Proof: The proof follows from the fact that as $\underline{\mathbf{C}}_N = \sigma_N^2 \mathbf{I}$, all eigenvalues $\mu_1 = \dots = \mu_N = \sigma_N^2$. ■

As mentioned before, due to dimensionality reduction the absolute information loss $L(\mathbf{X} \rightarrow \mathbf{Y}_M)$ is infinite; a direct consequence is that the irrelevant information loss, $L_{\mathbf{X}|\mathbf{S}}(\mathbf{X} \rightarrow \mathbf{Y}_M)$, is infinite as well. Given the assumptions of Corollary 3 holds, the PCA is a good solution to our signal enhancement problem.

VI. EXAMPLES

We now try to illustrate a few of our results with the help of some examples. In particular, we analyze a communication channel with binary input and a dimensionality-reducing PCA for non-Gaussian data. Unless otherwise noted, \log denotes the binary logarithm.

A. Communication channel with uniform noise

Consider the digital communication system depicted in Fig. 4. Let the data symbols S be uniformly distributed on $\{-1, 1\}$, thus $H(S) = 1$. Let further N be a noise signal which, for the sake of simplicity, is uniformly distributed on $[-a, a]$, $a > 1$. Clearly, $X = S + N$ is a continuous-valued signal with infinite entropy. During quantization an infinite amount of information is lost, $L(X \rightarrow Y) = \infty$ [1]. The

main point here is that most of the information at the input of the quantizer is information about the noise signal N . We will now make this statement precise.

With the differential entropy of X given by

$$h(X) = \frac{1}{a} \log 4a + \frac{a-1}{a} \log 2a \quad (36)$$

and with $h(X|S) = \log 2a$ we get the amount of relevant information available at the input of the quantizer:

$$I(X; S) = \frac{1}{a} \quad (37)$$

Choosing now a single quantizer threshold $\gamma_1 = 0$, i.e., $Y = \text{sgn}(X)$ we obtain a binary symmetric channel with cross-over probability $P_e = \frac{a-1}{2a}$. With $H_2(\cdot)$ being the binary entropy function the mutual information thus computes to $I(Y; S) = 1 - H_2(P_e)$ (see, e.g., [10]) and we obtain a relevant information loss of

$$L_S(X \rightarrow Y) = H_2\left(\frac{a-1}{2a}\right) - \frac{a-1}{a}. \quad (38)$$

Conversely, if we use two quantizer thresholds $\gamma_1 = 1 - a$ and $\gamma_2 = a - 1$ we obtain a ternary output RV Y . Interpreting any value $\gamma_1 \leq X \leq \gamma_2$ as an erasure, we obtain a binary erasure channel with erasure probability $\frac{a-1}{a}$. The corresponding mutual information is computed as $I(Y; S) = \frac{1}{a}$ (see, e.g., [10]), and the relevant information loss vanishes.

As this example shows, while our system actually destroys an infinite amount of information (to be precise, exactly 100% of the available information [2]), the relevant information loss can still be zero. Signal enhancement, in the sense of removing irrelevant information, was thus successful.

In this example it is tempting to identify the noise variable N with the irrelevant information $X|S$. However, this not necessarily leads to a correct result, as we show next.

To this end, we substitute the quantizer in Fig. 4 by a magnitude device, i.e., $Y' = |X|$. By the fact that the probability density function of X has even symmetry, it can be shown that $L(X \rightarrow Y') = H(X|Y') = 1$ (e.g., [12]). One can further show that, since the marginal distribution of Y' coincides with the conditional distributions $Y|S = -1$ and $Y|S = 1$, the mutual information $I(Y'; S) = 0$. Thus, $L_S(X \rightarrow Y') = \frac{1}{a}$. As an immediate consequence, $L_{X|S}(X \rightarrow Y') = \frac{a-1}{a}$.

Let us now determine the information loss relevant w.r.t. N . By showing that $L_N(X \rightarrow Y') \neq L_{X|S}(X \rightarrow Y')$ we argue that noise and irrelevant information are not necessarily identical. Observe that

$$L_N(X \rightarrow Y') = I(X; N|Y') \quad (39)$$

$$= H(X|Y') - H(X|N, Y') \quad (40)$$

$$= 1 - H(S + N|N, Y') \quad (41)$$

$$= 1 - H(S|N, |S + N|). \quad (42)$$

Given we know N , S is uncertain only if $|S + N|$ yields the same value for both $S = 1$ and $S = -1$. In other words, we have to require that $|N - 1| = |N + 1|$. Squaring both sides translates this to requiring $N = -N$, which is fulfilled only

TABLE I
INFORMATION LOSS (RELEVANT AND IRRELEVANT) FOR THE
COMMUNICATION CHANNEL WITH UNIFORM NOISE.

Loss	$Y = \text{sgn}(X)$	$Y = X $
$L(X \rightarrow Y)$	∞	1
$L_S(X \rightarrow Y)$	$H_2(\frac{a-1}{2a}) - \frac{a-1}{a}$	$\frac{1}{a}$
$L_{X S}(X \rightarrow Y)$	∞	$\frac{a-1}{a}$
$L_N(X \rightarrow Y)$	∞	1
$L_{X N}(X \rightarrow Y)$	$\frac{a-1}{2a}$	0

for $N = 0$. Since $\Pr(N = 0) = 0$ (N is a continuous RV), we automatically obtain

$$L_N(X \rightarrow Y') = 1 \neq \frac{a-1}{a} = L_{X|S}(X \rightarrow Y'). \quad (43)$$

Indeed, the reason why we cannot identify the noise with the irrelevant information can be related to N and S not being conditionally independent given X . We summarize these findings in Table I.

B. PCA with non-Gaussian data

Assume we observe two independent data sources, S_1 and S_2 , with three sensors which are corrupted by independent, unit-variance Gaussian noise N_i , $i = 1, 2, 3$. Our sensor signals shall be defined as

$$X_1 = S_1 + N_1, \quad (44)$$

$$X_2 = S_1 + S_2 + N_2, \text{ and} \quad (45)$$

$$X_3 = S_2 + N_3. \quad (46)$$

We assume further that our data sources have variances σ_1^2 and σ_2^2 and are non-Gaussian, but that they still can be described by a (joint) probability density function. We obtain the covariance matrix of $\mathbf{X} = [X_1, X_2, X_3]^T$ as

$$\underline{\mathbf{C}}_{\mathbf{X}} = \begin{bmatrix} \sigma_1 + 1 & \sigma_1 & 0 \\ \sigma_1 & \sigma_1 + \sigma_2 + 1 & \sigma_2 \\ 0 & \sigma_2 & \sigma_2 + 1 \end{bmatrix}. \quad (47)$$

Performing the eigenvalue decomposition yields three eigenvalues,

$$\{\lambda_1, \lambda_2, \lambda_3\} = \{\sigma_1 + \sigma_2 + 1 + C, \sigma_1 + \sigma_2 + 1 - C, 1\} \quad (48)$$

where $C = \sqrt{\sigma_1^2 + \sigma_2^2 - \sigma_1\sigma_2}$.

We now reduce the dimension of the output vector \mathbf{Y} from $N = 3$ to $M = 2$ by dropping the component corresponding to the smallest eigenvalue. Using Lemma 1 we now give an upper bound on the relevant loss, $L_S(\mathbf{X} \rightarrow \mathbf{Y}_M)$. To this end, note that \mathbf{N} is iid Gaussian, and thus $h(\tilde{\mathbf{N}}_c | \tilde{\mathbf{N}}_M) = h(\tilde{\mathbf{N}}_c)$. By assumption, $\underline{\mathbf{C}}_{\mathbf{N}} = \mathbf{I}$, and by the orthogonality of the transform,

$$h(\tilde{\mathbf{N}}_c) = \frac{1}{2} \ln(2\pi e). \quad (49)$$

Conversely, we know that

$$h(\mathbf{Y}_c | \mathbf{Y}_M) \leq h(\mathbf{Y}_{c,G} | \mathbf{Y}_{M,G}) = h(\mathbf{Y}_{c,G}) = \frac{1}{2} \ln(2\pi e \lambda_3). \quad (50)$$

With $\lambda_3 = 1$ and Lemma 1 we thus get

$$L_S(\mathbf{X} \rightarrow \mathbf{Y}_M) \leq 0. \quad (51)$$

The relevant information loss vanishes. Indeed, the eigenvector corresponding to the smallest eigenvalue is given as $\mathbf{p}_3 = \frac{1}{\sqrt{3}}[1, -1, 1]^T$; thus, for Y_3 one would obtain

$$Y_3 = \frac{X_1 + X_3 - X_2}{\sqrt{3}} = \frac{N_1 + N_3 - N_2}{\sqrt{3}}. \quad (52)$$

Since this component does not contain any signal portion, dropping it does not lead to a loss of relevant information.

Note that in case the noise sources N_i do not have the same variances, the application of PCA may lead to a loss of information, even though the relevant information is still concentrated on a subspace of lower dimensionality. We make this precise the next example.

C. PCA with large noise variances

We now again assume that we use three sensors to observe two independent, non-Gaussian data sources which are corrupted by independent Gaussian noise:

$$X_1 = S_1 + N_1, \quad (53)$$

$$X_2 = S_2 + N_2, \text{ and} \quad (54)$$

$$X_3 = N_3 \quad (55)$$

This time, however, we assume that the data sources have unit variance, and that the variance of noise source N_i is i , thus $\{\mu_1, \mu_2, \mu_3\} = \{3, 2, 1\}$. With this, the covariance matrix of \mathbf{X} is given as

$$\underline{\mathbf{C}}_{\mathbf{X}} = \begin{bmatrix} 1+1 & 0 & 0 \\ 0 & 1+2 & 0 \\ 0 & 0 & 3 \end{bmatrix} = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 3 \end{bmatrix} \quad (56)$$

and has eigenvalues $\{\lambda_1, \lambda_2, \lambda_3\} = \{3, 3, 2\}$.

Since $\underline{\mathbf{C}}_{\mathbf{X}}$ is already diagonal, the PCA only leads to an ordering w.r.t. the eigenvalues; dropping the component of \mathbf{Y} corresponding to the smallest eigenvalue, i.e., dimensionality reduction from $N = 3$ to $M = 2$, yields $Y_1 = N_3$ and $Y_2 = S_2 + N_2$. Since we do not have access to S_1 anymore, information is lost – in fact, PCA suggested to drop the most informative component of \mathbf{X} (the one with the highest SNR).

With Lemma 1 we can now compute the Gaussian upper bound on the relevant information loss: Following the proof of Theorem 1,

$$L_S(\mathbf{X} \rightarrow \mathbf{Y}_M) \leq h(Y_{3,G}) - h(N_1) = \frac{1}{2} \ln 2. \quad (57)$$

The bound obtained from Theorem 2 is loose here, evaluating to $L_S(\mathbf{X} \rightarrow \mathbf{Y}_M) < \frac{1}{2} \ln 3$.

Since in this case the noise is not iid, the condition of Theorem 1 is not fulfilled, thus the Gaussian upper bound is not minimized by the PCA. Indeed, by preserving X_1 and X_2 (and dropping X_3) the relevant information loss vanishes.

VII. CONCLUSION

In this work we presented the notion of relevant information loss and analyzed many of its elementary properties. We argued that relevant information loss is a central quantity in the problem of signal enhancement and thus, of system theory in general. A comparison with the literature about machine learning and neural networks revealed that many of the algorithms introduced there can be reformulated as solutions for the signal enhancement problem using the notion of relevant information loss. As an example, we discussed principle component analysis used for dimensionality reduction and derived conditions under which the relevant information loss vanishes.

Future work will concentrate on different blocks of the signal processing chain, such as quantizers, sampling devices, and filters as well as on investigating a possible connection between relevant information loss and state space aggregation.

APPENDIX

A. Proof of Proposition 2

The proof follows along the same lines as the proof of the data processing inequality for mutual information (see [10]). We first expand $I(X; W, V|Y)$ as follows:

$$I(X; W, V|Y) = I(X; V|Y) + I(X; W|V, Y) \quad (58)$$

$$= I(X; W|Y) + I(X; V|W, Y) \quad (59)$$

$$= I(X; W|Y) \quad (60)$$

where the last line follows from the fact that V and X are conditionally independent given W .

Using now Definition 1 we obtain

$$L_W(X \rightarrow Y) = I(X; W|Y) \quad (61)$$

$$= I(X; V|Y) + I(X; W|V, Y) \quad (62)$$

$$= L_V(X \rightarrow Y) + I(X; W|V, Y) \quad (63)$$

$$\geq L_V(X \rightarrow Y) \quad (64)$$

which completes the proof. ■

B. Proof of Corollary 1

If $f(S) - S - X - Y$ is a Markov chain, the proof follows immediately from Proposition 2. In case $S - f(S) - X - Y$, one gets with the proof of Proposition 2

$$L_{f(S)}(X \rightarrow Y) = L_S(X \rightarrow Y) + I(X; f(S)|S, Y) \quad (65)$$

$$= L_S(X \rightarrow Y). \quad (66)$$

Thus, in this case Proposition 2 is shown to hold with equality.

Furthermore, using Definition 1, the Corollary follows immediately from [15, Thm. 3.7.1]. ■

C. Proof of Proposition 4

For the proof, note that with Definition 1 we get

$$L_S(X \rightarrow Z) \stackrel{(a)}{=} I(S; X, Y|Z) \quad (67)$$

$$\stackrel{(b)}{=} I(S; X|Y, Z) + I(S; Y|Z) \quad (68)$$

$$\stackrel{(a)}{=} I(S; X|Y) + I(S; Y|Z) \quad (69)$$

where (a) is due to the fact that $Y = g(X)$ and $Z = h(Y)$, respectively, and (b) is the chain rule of information. ■

D. Proof of Lemma 1

We start by noting that $L_S(\mathbf{X} \rightarrow \mathbf{Y}_M) = L_S(\mathbf{Y} \rightarrow \mathbf{Y}_M)$ since \mathbf{Y} and \mathbf{X} are related by an invertible transform. Thus,

$$L_S(\mathbf{Y} \rightarrow \mathbf{Y}_M) = I(\mathbf{Y}; \mathbf{S}) - I(\mathbf{Y}_M; \mathbf{S}) \quad (70)$$

$$\begin{aligned} &= h(\mathbf{Y}) - h(\mathbf{Y}|\mathbf{S}) - h(\mathbf{Y}_M) + h(\mathbf{Y}_M|\mathbf{S}) \\ &= h(\mathbf{Y}_c|\mathbf{Y}_M) - h(\mathbf{W}^T \mathbf{S} + \mathbf{W}^T \mathbf{N}|\mathbf{S}) \\ &\quad + h(\mathbf{I}_M \mathbf{W}^T \mathbf{S} + \mathbf{I}_M \mathbf{W}^T \mathbf{N}|\mathbf{S}) \end{aligned} \quad (71)$$

$$\begin{aligned} &\stackrel{(a)}{=} h(\mathbf{Y}_c|\mathbf{Y}_M) - h(\mathbf{W}^T \mathbf{N}) + h(\mathbf{I}_M \mathbf{W}^T \mathbf{N}) \\ &= h(\mathbf{Y}_c|\mathbf{Y}_M) - h(\tilde{\mathbf{N}}_c|\tilde{\mathbf{N}}_M) \end{aligned} \quad (72)$$

where (a) is due to independence of \mathbf{S} and \mathbf{N} . This completes the proof. ■

E. Proof of Theorem 1

By Lemma 1 and Definition 2,

$$L_S(\mathbf{X} \rightarrow \mathbf{Y}_M) = h(\mathbf{Y}_c|\mathbf{Y}_M) - h(\tilde{\mathbf{N}}_c|\tilde{\mathbf{N}}_M) \quad (73)$$

$$\begin{aligned} &= h(\mathbf{Y}_{c,G}|\mathbf{Y}_{M,G}) - J(\mathbf{Y}_c|\mathbf{Y}_M) \\ &\quad - h(\tilde{\mathbf{N}}_{c,G}|\tilde{\mathbf{N}}_{M,G}) + J(\tilde{\mathbf{N}}_c|\tilde{\mathbf{N}}_M) \end{aligned} \quad (74)$$

$$\leq h(\mathbf{Y}_{c,G}|\mathbf{Y}_{M,G}) - h(\tilde{\mathbf{N}}_{c,G}|\tilde{\mathbf{N}}_{M,G}) \quad (75)$$

$$\stackrel{(a)}{=} h(\mathbf{Y}_{c,G}) - h(\tilde{\mathbf{N}}_{c,G}|\tilde{\mathbf{N}}_{M,G}) \quad (76)$$

$$\stackrel{(b)}{=} h(\mathbf{Y}_{c,G}) - h(\tilde{\mathbf{N}}_{c,G}) \quad (77)$$

$$= \frac{1}{2} \ln \left(\prod_{i=M+1}^N \frac{\lambda_i}{\mu} \right). \quad (78)$$

Here, (a) is due to the fact that the PCA decorrelates the output data \mathbf{Y} and thus leads to independence of $\mathbf{Y}_{c,G}$ and $\mathbf{Y}_{M,G}$ (in the sense of Definition 2). By similar reasons (b) follows from the fact that \mathbf{N} is iid ($\mathbf{C}_\mathbf{N}$ is a scaled identity matrix, with all eigenvalues being equal $\mu_i = \mu$). Since the PCA preserves the M coordinates corresponding to the M largest eigenvalues of $\mathbf{C}_\mathbf{Y}$, the last line (obtained with [10, Thm. 8.4.1] and [16, Fact 5.10.14]) represents the smallest Gaussian upper bound and completes the proof. ■

F. Proof of Theorem 2

We recapitulate (76) from the proof of Theorem 1 and get

$$L_S(\mathbf{X} \rightarrow \mathbf{Y}_M) \leq h(\mathbf{Y}_{c,G}) - h(\tilde{\mathbf{N}}_{c,G}|\tilde{\mathbf{N}}_{M,G}) \quad (79)$$

$$= h(\mathbf{Y}_{c,G}) - h(\tilde{\mathbf{N}}_G) + h(\tilde{\mathbf{N}}_{M,G}). \quad (80)$$

With [10, Thm. 8.4.1] and [16, Fact 5.10.14] we get

$$h(\tilde{\mathbf{N}}_G) = \frac{1}{2} \ln \left((2\pi e)^N \prod_{i=1}^N \mu_i \right) \quad (81)$$

and

$$h(\mathbf{Y}_{c,G}) = \frac{1}{2} \ln \left((2\pi e)^{N-M} \prod_{i=M+1}^N \lambda_i \right). \quad (82)$$

If now $\mathbf{C}_{\tilde{\mathbf{N}}}$ denotes the $(M \times M)$ -covariance matrix of $\tilde{\mathbf{N}}_M$ (and, thus, of $\tilde{\mathbf{N}}_{M,G}$) and $\{\tilde{\mu}_i\}$ the set of eigenvalues of $\mathbf{C}_{\tilde{\mathbf{N}}}$,

we obtain

$$L_S(\mathbf{X} \rightarrow \mathbf{Y}_M) \leq \frac{1}{2} \ln \left(\frac{\prod_{i=M+1}^N \lambda_i \prod_{i=1}^M \tilde{\mu}_i}{\prod_{i=1}^N \mu_i} \right). \quad (83)$$

We now complete the proof by providing upper bounds on the eigenvalues in the numerator. It is easy to verify that $\underline{\mathbf{C}}_{\tilde{N}}$ is the top left principal submatrix of $\underline{\mathbf{W}}^T \underline{\mathbf{C}}_N \underline{\mathbf{W}}$ (which, by the orthogonality of $\underline{\mathbf{W}}$ has the same eigenvalues as $\underline{\mathbf{C}}_N$). As a consequence, we can employ Cauchy's interlacing inequality [16, Thm. 8.4.5]:

$$\mu_{i+N-M} \leq \tilde{\mu}_i \leq \mu_i \quad (84)$$

The second bound, $\lambda_i \leq \mu_1$, is derived from Weyl's inequality [16, Thm. 8.4.11]

$$\lambda_i \leq \nu_i + \mu_1 \quad (85)$$

and by noticing that $\nu_j = 0$ for all $j > M$.

Combining all this, we obtain as an upper bound on the information loss

$$L_S(\mathbf{X} \rightarrow \mathbf{Y}_M) \leq \frac{1}{2} \ln \left(\frac{\prod_{i=M+1}^N \lambda_i \prod_{i=1}^M \tilde{\mu}_i}{\prod_{i=1}^N \mu_i} \right) \quad (86)$$

$$\leq \frac{1}{2} \ln \left(\frac{\prod_{i=M+1}^N \mu_1 \prod_{i=1}^M \tilde{\mu}_i}{\prod_{i=1}^N \mu_i} \right) \quad (87)$$

$$= \frac{1}{2} \ln \left(\prod_{i=M+1}^N \frac{\mu_1}{\mu_i} \right). \quad (88)$$

This completes the proof. ■

REFERENCES

- [1] B. C. Geiger and G. Kubin, "On the information loss in memoryless systems: The multivariate case," in *Proc. Int. Zurich Seminar on Communications (IZS)*, Zurich, Feb. 2012, pp. 32–35, extended version available: [arXiv:1109.4856 \[cs.IT\]](#).
- [2] —, "Relative information loss in the PCA," in *Proc. IEEE Information Theory Workshop (ITW)*, Lausanne, Sep. 2012, pp. 562–566, extended version available: [arXiv:1204.0429 \[cs.IT\]](#).
- [3] D. H. Johnson, "Information theory and neural information processing," *IEEE Trans. Inf. Theory*, vol. 56, no. 2, pp. 653–666, Feb. 2010.
- [4] M. Plumbley, "Information theory and unsupervised neural networks," Cambridge University Engineering Department, Tech. Rep. CUED/F-INFENG/TR. 78, 1991.
- [5] G. Deco and D. Obradovic, *An Information-Theoretic Approach to Neural Computing*. New York, NY: Springer, 1996.
- [6] N. Tishby, F. C. Pereira, and W. Bialek, "The information bottleneck method," in *Proc. Allerton Conf. on Communication, Control, and Computing*, Sep. 1999, pp. 368–377.
- [7] M. A. Sánchez-Montañés and F. J. Corbacho, "A new information processing measure for adaptive complex systems," *IEEE Trans. Neural Netw.*, vol. 15, no. 4, pp. 917–927, Jul. 2004.
- [8] M. A. Sánchez-Montañés, "A theory of information processing for adaptive systems: Inspiration from biology, formal analysis and application to artificial systems," Ph.D. dissertation, Universidad Autónoma de Madrid, Jun. 2003.
- [9] J. C. Principe, *Information Theoretic Learning: Renyi's Entropy and Kernel Perspectives*, ser. Information Science and Statistics. New York, NY: Springer, 2010.
- [10] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. Hoboken, NJ: Wiley Interscience, 2006.
- [11] R. Linsker, "Self-organization in a perceptual network," *IEEE Computer*, vol. 21, no. 3, pp. 105–117, Mar. 1988.
- [12] B. C. Geiger, C. Feldbauer, and G. Kubin, "Information loss in static nonlinearities," in *Proc. IEEE Int. Sym. Wireless Communication Systems (ISWSC)*, Aachen, Nov. 2011, pp. 799–803, extended version available: [arXiv:1102.4794 \[cs.IT\]](#).
- [13] N. Slonim and N. Tishby, "Agglomerative information bottleneck," in *Advances in Neural Information Processing Systems (NIPS)*. MIT Press, 1999, pp. 617–623.
- [14] G. Chechik and N. Tishby, "Extracting relevant structures with side information," in *Advances in Neural Information Processing Systems (NIPS)*. MIT press, 2002, pp. 857–864. [Online]. Available: [books.nips.cc/papers/files/nips15/AA48.pdf](#)
- [15] M. S. Pinsker, *Information and Information Stability of Random Variables and Processes*. San Francisco, CA: Holden Day, 1964.
- [16] D. S. Bernstein, *Matrix Mathematics*. Princeton, NJ: Princeton University Press, 2005.